

Analyzing Adaptive Cache Replacement Strategies

Leo Shao¹, Mario E. Consuegra², Raju Rangaswami¹, and Giri Narasimhan¹

¹ School of Computing and Information Sciences, Florida International University, Miami, FL 33199, USA. {lshao002,raju,giri}@fiu.edu

² Microsoft Corporation, Seattle, WA, USA.

Abstract

Adaptive Replacement Cache (ARC) and *CLOCK with Adaptive Replacement (CAR)* are state-of-the-art “adaptive” cache replacement algorithms invented to improve on the shortcomings of classical cache replacement policies such as *LRU* and *CLOCK*. Both *ARC* and *CAR* have been shown to outperform their classical and popular counterparts in practice. However, for over a decade, no theoretical proof of the competitiveness of *ARC* and *CAR* is known. We prove that for a cache of size N , (a) *ARC* is $4N$ -competitive, and (b) *CAR* is $21N$ -competitive, thus proving that no “pathological” worst-case request sequence exists that could make them perform much worse than *LRU*.

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems

Keywords and phrases Cache replacement, LRU, ARC, CLOCK, CAR

1 Introduction

In a review by Irani and Karlin (see [5], Chapter 13), they state that although online problems are pervasive, the **cache replacement problem** (they refer to it as the “paging problem” in a two-level store) *is the most fundamental and practically important online problem in computer science*. Megiddo and Modha [12] call caching a “fundamental metaphor in modern computing”.

In a seminal paper, Sleator and Tarjan [14] introduced the concept of **competitive analysis** to analyze the performance of *online algorithms* by comparing them against the performance of an *optimal offline* algorithm (**OPT**). They used their model to analyze online algorithms such as *cache replacement algorithms*, proving two important results in the process. First, they showed that under the “demand paging” model (no prefetching), no online algorithm for cache replacement could achieve a competitiveness ratio less than N , where N is the size of the cache. Second, they showed that the *Least Recently Used* (**LRU**) cache replacement scheme has a competitiveness ratio of N , suggesting that under this measure, **LRU** is optimal. These results are significant. However, we note an important, but subtle point that highlights the difference between theory and practice. Sleator and Tarjan’s techniques analyze online algorithms in terms of their *worst-case* behavior (i.e., over all possible inputs), which means that other algorithms with poorer competitiveness ratios could perform better in practice. Another way to state this is that the results assume an *oblivious adversary* who designs the inputs for the online algorithms in a way that make them perform as poorly as possible. The upper bound on performance ratio merely guarantees that no surprises are in store, i.e., there is no input designed by an adversary that can make the algorithm perform poorly.

The **LRU** algorithm was considered the most optimal page replacement policy for a long time, but it had the drawback of not being “scan-resistant”, i.e., items used only once could pollute the cache and diminish its performance. Furthermore, **LRU** is difficult to implement efficiently, since moving an accessed item to the front of the queue is an expensive operation, first requiring locating the item, and then requiring data moves that could lead to unacceptable cache contention if it is to be implemented consistently and correctly. The **CLOCK** algorithm was invented by Frank Corbató in 1968 as an efficient one-bit approximation to **LRU** with minimum overhead [3] and continues to be used in MVS, Unix, Linux, and Windows operating systems [4]. Like **LRU**, **CLOCK** is also not scan-resistant because it puts too much emphasis on “recency” of access and pays no attention to “frequency” of access. So there are sequences in which many other algorithms can have significantly less cost than the theoretically optimal **LRU**. Since then, many other cache replacement strategies have been developed and have been showed to be better than **LRU** in practice. These are discussed below in Section 2.

An important development in this area was the invention of **adaptive algorithms**. While regular “online” algorithms are usually designed to respond to input requests in an optimal manner, these *self-tuning* algorithms are capable of adapting to changes in the request sequence caused by changes in the workloads. It adapts to increase or decrease in the number of items that have been accessed more than once. Modha and Megiddo [11] developed a self-tuning algorithm called *Adaptive Replacement Strategy* (**ARC**), a hybrid of **LFU** and **LRU**. Soon afterwards, another algorithm called **CAR** was also developed. **CAR** is a hybrid of **LFU** and **CLOCK** [2]. It had the adaptive flavor of **ARC**, but was more efficient. Experiments show that **ARC** and **CAR** outperform **LRU** and **CLOCK** for many benchmark data sets [2]. Versions of **ARC** have been deployed in commercial systems such as the IBM

DS6000/DS8000, Sun Microsystems’s ZFS, and in PostgreSQL.

Unfortunately, no **theoretical analysis** of the various adaptive algorithms exists in the literature. The main open question that remained unanswered was whether or not there exists some “pathological” request sequence that could force **ARC** or **CAR** to perform poorly. We settle this open question by showing that **ARC** is $4N$ -competitive and **CAR** is $21N$ -competitive, thus proving that under Sleator and Tarjan’s theoretical model, their worst-case behavior is asymptotically optimal, and therefore not much worse than the theoretically optimal **LRU**. The main challenge in solving these problems is that of carefully designing the potential function for the analysis.

After providing relevant background on cache replacement algorithms in Section 2, we discuss the competitive analysis of **ARC** in Section 3 and the competitive analyses of **CLOCK** and **CAR** in Section 4. Concluding remarks can be found in Section 5.

2 Previous Work on Cache Replacement Strategies

In practice, since **LRU** evicts the least recently used entry, it tends to perform well when there are many items that are requested more than once in a relatively short period of time, and performs poorly on “scans”. **CLOCK** implements a circular buffer for the entries, and its replacement strategy involves cycling through the pages in that buffer, treating it like a clock. Instead of the time of reference as in **LRU**, **CLOCK** maintains a reference bit for each item. If an entry is requested, the reference bit is set. Each entry is considered for eviction when the clock hand points to it. It is evicted only if its reference bit is not set, else the reference bit is reset and the clock moves on to the next entry.

The **DUELINGCLOCK** algorithm [6] is like **CLOCK** but keeps the clock hand at the newest page rather than the oldest one, which allows it to be scan-resistant. More recent algorithms try to improve over **LRU** by implementing multiple cache levels and leveraging history. In [13] the **LRU- K** algorithm was introduced. Briefly, the **LRU- K** algorithm estimates interarrival times from observed requests, and favors retaining pages with shorter interarrival times. Experiments have shown **LRU-2** performs better than **LRU**, and that **LRU- K** does not show increase in performance over **LRU-2** [13], but has a higher implementation overhead. It was also argued that **LRU- K** is optimal under the independence reference model (IRM) among all algorithms A that have limited knowledge of the K most recent references to a page and no knowledge of the future [13].

In essence, the **LRU- K** algorithm tries to efficiently approximate *Least Frequently Used* (**LFU**) cache replacement algorithm. As K becomes larger, it gets closer and closer to **LFU**. It has been argued that **LFU** cannot adapt well to changing workloads because it may replace currently “hot” blocks instead of “cold” blocks that had been “hot” in the past. **LFU** is implemented as a heap and takes $O(\log N)$ time per request.

Another cache replacement algorithm is **LIRS** [8]. The **LIRS** algorithm evicts the page with the largest IRR (inter-reference recency). It attempts to keep a small ($\approx 1\%$) portion of the cache for HIR (high inter-reference) pages, and a large ($\approx 99\%$) portion of the cache for LIR (low inter-reference) pages. The **CLOCK-PRO** algorithm approximates **LIRS** efficiently using **CLOCK** [7]. The **2Q** [9] algorithm is scan-resistant. It keeps a FIFO buffer A_1 of pages that have been accessed once and a main **LRU** buffer A_m of pages accessed more than once. **2Q** admits only hot pages to the main buffer. The buffer A_1 is divided into a main component that keeps the pages in A_1 that still reside in cache, and a history component that remembers pages that have been evicted after one access. The relative sizes of the main and history components are tunable parameters. **2Q** has time complexity of $O(1)$. Another

algorithm that tries to bridge the gap between recency and frequency is **LRFU** [10]. This is a hybrid of **LRU** and **LFU** and is adaptive to changes in workload. The time complexity ranges from $O(1)$ for **LRU** to $O(\log n)$ for **LFU**.

3 Analyzing ARC

To facilitate our discussion, we briefly describe the **ARC** algorithm. As mentioned before, it combines ideas of recency and frequency. **ARC**'s cache is organized into a “main” part (of size N) and a “history” part (of size N). The main part is further divided into two lists, T_1 and T_2 , both sorted by “recency”. T_1 focuses on “recency” because it contains pages with short-term utility. Consequently, when an item is accessed for the first time from the disk, it is brought into T_1 . Items “graduate” to T_2 when they are accessed more than once. Thus, T_2 deals with “frequency” and stores items with long-term utility. Additionally, **ARC** maintains a history of N more items, consisting of B_1 , i.e., items that have been recently deleted from T_1 , and B_2 , i.e., items that have been recently deleted from T_2 . History lists are also organized in the order of recency of access. The unique feature of **ARC** is its self-tuning capability, which makes it scan-resistant. Based on a self-tuning parameter, p , the size of T_1 may grow or shrink relative to the size of T_2 . The details of the algorithm are fairly complex and non-intuitive. Detailed pseudocode for **ARC** (Figure 4 from [11]) is provided in the Appendix for convenience.

It must be noted that we make one minor change in the **ARC** algorithm as presented by Megiddo and Modha. In the original algorithm, when an item is missing from the cache, but is found in the history, then the adaptive parameter p increases or decreases depending on whether it was found in B_1 or B_2 . The change is assumed to be by a quantity δ , given by $+\max\{|B_2|/|B_1|, 1\}$ or $-\max\{|B_1|/|B_2|, 1\}$. For our analysis, we replace the change by the quantities $+1$ and -1 respectively. At this point, it is unclear if our analysis extends to the original

3.1 Definitions and Notation

We start with some notation and definitions. If X is the set of pages in a cache, then let $MRU(X)$ and $LRU(X)$ be the most recently and least recently used pages from X . Let $MRU_k(X)$ and $LRU_k(X)$ be the k most recently and k least recently used pages from X .

Let lists L_1 (and L_2) be the lists obtained by concatenating lists T_1 and B_1 (T_2 and B_2 , resp.). We let $\ell_1, \ell_2, t_1, t_2, b_1, b_2$ denote the sizes of $L_1, L_2, T_1, T_2, B_1, B_2$, respectively. Finally, let $t := t_1 + t_2$ and $\ell := \ell_1 + \ell_2$.

In what follows, we will assume that the two algorithms **OPT** and **ARC** are run in parallel, each maintaining their own caches as per their own replacement policies. At any instant of time during the parallel simulation of **OPT** and **ARC**, and for any list X , we let $MRU_k(X)$ be denoted by $TOP(X)$, where k is the largest integer such that all pages of $MRU_k(X)$ are also in the cache maintained by **OPT**. We let L'_1, L'_2, T'_1, T'_2 denote the TOP s of L_1, L_2, T_1, T_2 , respectively, with sizes $\ell'_1, \ell'_2, t'_1, t'_2$, respectively. We let b'_1 and b'_2 denote the sizes of the $B'_1 = L'_1 \cap B_1$ and $B'_2 = L'_2 \cap B_2$, respectively. Note that if $b'_1 > 0$ ($b'_2 > 0$, resp.), then all of T_1 (T_2 , resp.) is in **OPT**. Finally, we let $\ell' := \ell'_1 + \ell'_2$.

For all our proofs, we assume that algorithm X being analyzed is provided an arbitrary request sequence $\sigma = \sigma_1, \sigma_2, \dots, \sigma_m$. Following the strategy outlined by Albers [1], we partition σ into phases $P(0), P(1), \dots$, such that X has exactly N faults on $P(i)$, for every $i \geq 0$. Also, let the subsequence of the request sequence σ in phase $P(i)$ be $\sigma(i)$.

3.2 Analyzing the Competitiveness of ARC

Let C_{arc} and C_{opt} be the costs incurred by the algorithms **ARC** and **OPT**. Below, we prove the competitiveness of **ARC**. In any given phase $P(i)$, if the faults are to N distinct pages, and if it is different from the last fault prior to the start of phase $P(i)$, then **OPT** is guaranteed to have a fault, allowing for the competitiveness of **ARC** to be easily proved. However, since this is not guaranteed, a more complicated proof is warranted.

We define the potential function as follows:

$$\Phi = p - [(b'_1 - t) + 2 \cdot (t'_1 - t) + 3 \cdot (b'_2 - t) + 4 \cdot (t'_2 - t)] \quad (1)$$

The main result of this section is the following theorem:

► **Theorem 1.** *Algorithm ARC is $4N$ -competitive.*

We will prove the above theorem by proving the following inequality for any phase $P(i)$, $i \geq 0$. Note that ΔX represents the change in quantity X .

$$C_{arc}(\sigma(i)) + \Delta\Phi \leq 4N \cdot C_{opt}(\sigma(i)) \quad (2)$$

Summing up the above inequality for all requests in all phases would prove the theorem. We first state an observation about phase $P(0)$ as a lemma.

► **Lemma 2.** *At the end of phase $P(0)$, **ARC**'s cache is "full", i.e., $t = N$. The cache remains full from that point onward.*

Proof. Note that in phase $P(0)$, the request sequence $\sigma(0)$ requests N distinct pages. At the end of phase $P(0)$, the cache maintained by **ARC** is full, i.e., $t = N$. Since evictions only happen when **ARC** has a miss, its cache will remain full from this point onward. ◀

In phase $P(0)$, if **ARC** and **OPT** start with the same cache contents, then the first time **ARC** faults in this phase, so will **OPT**, thereby guaranteeing that **OPT** faults at least once during this phase. Since **ARC** faults exactly N times during this phase, inequality 2 holds for phase $P(0)$.

To analyze phase $P(i)$, $i > 0$, we will prove that for every individual request, $\sigma_j \in \sigma(i)$:

$$C_{arc}(\sigma_j) + \Delta\Phi \leq 4N \cdot C_{opt}(\sigma_j) \quad (3)$$

Summing up the above inequality for all requests, $\sigma_j \in \sigma(i)$ will prove inequality (2) for phase $P(i)$.

We assume that request σ_j is processed in two distinct steps: first when **OPT** services the page request and, next when **ARC** services the request. We will show that inequality (3) is satisfied for each of the two steps.

Step 1: OPT serves request σ_j .

Since only **OPT** acts in this step, $C_{ARC} = 0$, and t does not change. There are two possible cases: either **OPT** faults on σ_j or it does not.

If **OPT** does not fault on this request, then $C_{OPT} = 0$. Since the contents of the cache maintained by **OPT** does not change, and neither do the quantities b'_1, b'_2, t'_1, t'_2 , we have $\Delta\Phi = 0$, and $C_{arc}(\sigma(i)) + \Delta\Phi \leq 4N \cdot C_{opt}(\sigma(i)) \leq 0$.

If **OPT** faults on request $\sigma(i)$, then $C_{OPT} = 1$. The contents of the cache maintained by **OPT** does change, which will affect the potential function. Since $0 \leq b'_1 + t'_1 + b'_2 + t'_2 \leq N$, the quantity $b'_1 + 2t'_1 + 3b'_2 + 4t'_2$ cannot change by more than $4N$. Thus, $C_{arc}(\sigma_j) + \Delta\Phi \leq 4N$, proving inequality (3).

Step 2: **ARC** serves request σ_j .

There are four possible cases. Case 1 deals with the case when **ARC** finds the page in its cache. The other three cases assume that **ARC** faults on this request because the item is not in $T_1 \cup T_2$. Cases 2 and 3 assume that the missing page is found recorded in the history in lists B_1 and B_2 , respectively. Case 4 assumes that the missing page is not recorded in history.

Case I: **ARC** has a page hit.

Clearly, $C_{arc} = 0$. We consider several subcases. In each case, the requested page will be moved to $MRU(T_2)$ while shifting other pages in T_2 down.

Case I.1 If the requested page is in T'_1 , the move of this page from T'_1 to T'_2 implies $\Delta t'_1 = -1$; $\Delta t'_2 = +1$ and $\Delta \Phi = -(2 \cdot \Delta t'_1 + 4 \cdot \Delta t'_2) = -2$.

Case I.2 If the requested page is in T'_2 , the move of this page to $MRU(T_2)$ does not change the set of items in T'_2 . Thus, $\Delta t'_1 = \Delta t'_2 = 0$ and $\Delta \Phi = 0$.

Case I.3 If the requested page is in $T_1 - T'_1$, it is similar to the case when the requested page is in T'_1 , except that it could cause t'_1 and b'_1 to increase by a large amount. This could occur if the requested item was immediately following T'_1 and its move caused items below it to become part of T'_1 . Thus we have, $\Delta t'_2 = +1$ and $\Delta t'_1, \Delta b'_1 \geq 0$. Finally,

$$\begin{aligned} \Delta \Phi &\leq -(\Delta b'_1 + 2\Delta t'_1 + 4\Delta t'_2) \\ &\leq -4 \end{aligned}$$

Case I.4 If the requested page is in $T_2 - T'_2$, it is similar to the case when the requested page is in T'_2 , except that it could cause t'_2 and b'_2 to increase by a large amount. Thus we have, $\Delta b'_2 \geq 0$ and $\Delta t'_2 \geq 1$. Finally,

$$\begin{aligned} \Delta \Phi &\leq -(3\Delta b'_2 + 4\Delta t'_2) \\ &\leq -4 \end{aligned}$$

Next we will analyze the three cases when the requested page is not in **ARC**'s cache. Since $C_{arc} = 1$, the change in potential must be at most -1 in order for inequality (3) to be satisfied. We make the following useful observations in the form of lemmas.

► **Lemma 3.** *If **ARC** has a miss and if the page is not in **ARC**'s history, we have $\ell' = t'_1 + t'_2 + b'_1 + b'_2 < N$. Consequently, we also have $\ell'_1 < N$ and $\ell'_2 < N$.*

Proof. Since **OPT** has just finished serving the request, the page is present in the cache maintained by **OPT** just before **ARC** starts to serve the request. If **ARC** has a miss, there is at least one page in the cache maintained by **OPT** that is not present in the cache maintained by **ARC**, implying that $\ell' < N$. By definition, $\ell' = \ell'_1 + \ell'_2 = t'_1 + t'_2 + b'_1 + b'_2$. Thus, the lemma holds. ◀

► **Lemma 4.** *A call to procedure REPLACE either causes an element to be moved from T_1 to B_1 or from T_2 to B_2 . In either case, the change in potential due to REPLACE, denoted by $\Delta \Phi_R$, has an upper bound of 1.*

Proof. Procedure REPLACE is only called when **ARC** has a page miss. Clearly, it causes an element to be moved from T_1 to B_1 or from T_2 to B_2 . If that element is in T'_1 or T'_2 , then either $T_1 = T'_1$ or $T_2 = T'_2$ and the moved element becomes part of B'_1 or B'_2 . Thus, the potential change is +1. If that element is in $T_1 - T'_1$ ($T_2 - T'_2$, resp.), then B'_1 (B'_2 , resp.) was empty before the move and remains empty after the move. Hence, there is no change in potential. ◀

► **Lemma 5.** *After phase $P(0)$, every eviction is the result of an **ARC** miss and involves either $LRU(L_1)$ or $LRU(L_2)$. However, this page will be neither from L'_1 nor L'_2 .*

Proof. By Lemma 2, we know that at the end of phase $P(0)$, the cache maintained by **ARC** is full, i.e., $t = t_1 + t_2 = N$. Every miss after that involves moving an item from T_1 to B_1 or from T_2 to B_2 , followed by a possible eviction from either L_1 or L_2 . The proof is by contradiction. Assume that a page is evicted from either $L_1 \cap L'_1$ or $L_2 \cap L'_2$.

First, assume that the evicted page is from $L_1 \cap L'_1$. The only time when **ARC** will remove a page from L_1 is when $\ell_1 = \ell'_1 = N$ (Case IV A from the **ARC** algorithm). By Lemma 3, this is impossible since when **ARC** has a miss, $\ell'_1 < N$.

Next, assume that the evicted page is from $L_2 \cap L'_2$, implying that $\ell_2 = \ell'_2$. The only time when **ARC** will remove a page from L_2 is when $l_1 < N$ (case IV B from **ARC** algorithm), implying that that $l_2 = l'_2 \geq N + 1$, which is impossible since **OPT** can have at most N items. \blacktriangleleft

► **Lemma 6.** *After phase $P(0)$, if $T_1 = T'_1$ then a page in T'_2 will not be moved to B_2 . Similarly, if $T_2 = T'_2$ then a page from T'_1 will not be moved to B_1 .*

Proof. In an attempt to prove by contradiction, let us assume that $T_1 = T'_1$ and $T_2 = T'_2$ are simultaneously true and **ARC** has a miss. By Lemma 2, the cache maintained by **ARC** is full after phase $P(0)$. Thus, we have $t = t_1 + t_2 = N = t'_1 + t'_2$, which is impossible by Lemma 3.

Thus, on an **ARC** miss, if $T_1 = T'_1$, then $T_2 \neq T'_2$. Thus if $LRU(T_2)$ is moved to B_2 , this item cannot be from T'_2 . By a symmetric argument, if $T'_2 = T_2$, then $T_1 \neq T'_1$, and $LRU(T_1)$ is not in T'_1 . \blacktriangleleft

Case II: **ARC** has a miss and the missing page is in B_1

Note that in this case the value of p will change by $+1$, unless its value equals N , in which case it has no change. Thus $\Delta p \leq 1$.

If the missing item is in B'_1 , then $\Delta b'_1 = -1$ and $\Delta t'_2 = +1$. Adding the change due to **REPLACE**, we get

$$\begin{aligned} \Delta \Phi &\leq 1 - (\Delta b'_1 + 4 \cdot \Delta t'_2) + \Delta \Phi_R \\ &\leq -1 \end{aligned}$$

If the missing item is in $B_1 - B'_1$, then, as argued before, b'_1 could increase by a large amount if this item immediately follows items in B'_1 and is also followed by items in **OPT**. Thus, we have $\Delta t'_2 = 1$ and $\Delta b'_1 \geq 0$. Thus, we have

$$\begin{aligned} \Delta \Phi &\leq 1 - (\Delta b'_1 + 4 \cdot \Delta t'_2) + \Delta \Phi_R \\ &\leq -2 \end{aligned}$$

Case III: **ARC** has a miss and the missing page is in B_2 .

Note that in this case the value of p will change by -1 , if its value was positive, otherwise it has no change. Thus $\Delta p \leq 0$.

If the requested item is in B'_2 , then $\Delta t'_2 = 1$, and $\Delta b'_2 = -1$. Thus, we have

$$\begin{aligned} \Delta \Phi &= \Delta p - (3 \cdot \Delta b'_2 + 4 \cdot \Delta t'_2) + \Delta \Phi_R \\ &\leq 0 \end{aligned}$$

But this is not good enough since we need the potential change to be at most -1 . When $\Delta p = -1$, then we get the required inequality $\Delta \Phi \leq -1$. Clearly, the difficulty is when $\Delta p = 0$. Since the missing item is from b'_2 , it implies that B'_2 is non-empty and $T'_2 = T_2$. By Lemma 6 above, since the cache is full (i.e., $t_1 + t_2 = N$), there must be at least one item in

$T_1 - T'_1$, which means that means that $t_1 > 0$. As per the algorithm, since T_1 is non-empty and $p = 0$, we are guaranteed to replace $LRU(T_1)$, and not an element from T'_1 . Therefore, REPLACE will leave t'_1 and b'_1 unchanged, implying that $\Delta\Phi_R = 0$. Thus, we have

$$\begin{aligned}\Delta\Phi &= \Delta p - (3 \cdot \Delta b'_2 + 4 \cdot \Delta t'_2) + \Delta\Phi_R \\ &\leq -1\end{aligned}$$

If the requested item is from $B_2 - B'_2$, then $\Delta t'_2 \geq 1$, and $\Delta b'_2 \geq 0$. Thus, we have

$$\begin{aligned}\Delta\Phi &\leq \Delta p - (3 \cdot \Delta b'_2 + 4 \cdot \Delta t'_2) + \Delta\Phi_R \\ &\leq -3\end{aligned}$$

Case IV: ARC has a miss and the missing page is not in $B_1 \cup B_2$

We consider two cases. First, when $\ell_1 = N$, ARC will evict the $LRU(L_1)$. Since by Lemma 3, $\ell'_1 < N$, we know that for this case, b'_1 remains 0 and $\Delta t'_1 = +1$. Thus, we have

$$\begin{aligned}\Delta\Phi &\leq -(2 \cdot \Delta t'_1) + \Delta\Phi_R \\ &\leq -1\end{aligned}$$

On the other hand, if $\ell_1 < N$, then ARC will evict the $LRU(L_2)$. Again, if the cache is full (i.e., $t_1 + t_2 = N$ and $\ell_1 + \ell_2 = 2N$), then we know that $\ell_2 > N$, which means that $L'_2 \neq L_2$ and $LRU(L_2)$ is not in L'_2 . Thus, deletion of $LRU(L_2) = LRU(B_2)$ will not affect b'_2 or any of the other quantities in the potential function. Then comes the REPLACE step, for which a bound has been proved earlier. Finally, a new item is brought in and placed in $MRU(T_1)$. Thus $\Delta t'_1 \leq 1$. Putting it all together, we have

$$\begin{aligned}\Delta\Phi &\leq -(2 \cdot \Delta t'_1) + \Delta\Phi_R \\ &\leq -1\end{aligned}$$

Wrapping up the proof of Theorem 1

Tying it all up, we have shown that inequality (3) holds for every request in the phase $P(i), i \geq 0$, proving that inequality (2) holds for every phase. We showed at the start that inequality (2) also holds for phase $P(0)$. Putting it all together, we have that

$$C_{arc}(\sigma) + \Delta\Phi \leq 4N \cdot C_{opt}(\sigma).$$

If we assume that the caches started empty, then the initial potential is 0, while the final potential can be at most $4N$. Thus, we have

$$C_{arc}(\sigma) \leq 4N \cdot C_{opt}(\sigma) + 4N,$$

thus proving Theorem 1. ◀

4 Analyzing CAR

As mentioned earlier, CLOCK with Adaptive Replacement (CAR) algorithm is a modification of the CLOCK algorithm and is inspired by the ARC algorithm. We start with a quick description of the CLOCK algorithm.

CLOCK was designed as a low-cost variant of LRU, requiring only one bit per item and one pointer to implement. CLOCK's cache is organized as a single "circular" list. The algorithm maintains a pointer to the "head" of the list. The item immediately counterclockwise to it

is the “tail” of the list. Each item is associated with a “mark” bit. Some of the pages in the cache are marked, and the rest are unmarked. When a page hit occurs that page is marked. The contents of the cache remain unchanged. When a page miss occurs, in order to make room for the requested page, the head page is evicted if the page is unmarked. If the head page is marked, the page is unmarked and the head is moved forward clockwise. After a page has been evicted, the requested page is unmarked and placed at the tail of the list.

Inspired by **ARC**, **CAR**’s cache is organized into two main lists, T_1 and T_2 , and two history lists, B_1 and B_2 . Inspired by **CLOCK**, both T_1 and T_2 are organized as “circular” lists, with each item associated with a mark bit. The history lists, B_1 and B_2 are maintained as simple FIFO lists. We let t_1, t_2, b_1, b_2 denote the sizes of T_1, T_2, B_1, B_2 , respectively. Also, let $t := t_1 + t_2$. Let lists L_1 (and L_2 , reps.) be the list obtained by concatenating list B_1 to the tail of T_1 (concatenating B_2 to the tail of T_2 , resp.), with sizes ℓ_1 (and ℓ_2 , resp.). We let T_1^0 and T_2^0 denote the unmarked pages in T_1 and T_2 , respectively; similarly, we let T_1^1 and T_2^1 denote the marked pages in T_1 and T_2 , respectively.

The following 7 invariants are maintained by **CAR** for the lists: (I1) $0 \leq t_1 + t_2 \leq N$; (I2) $0 \leq t_1 + b_1 \leq N$; (I3) $0 \leq t_2 + b_2 \leq 2N$; (I4) $0 \leq t_1 + t_2 + b_1 + b_2 \leq 2N$; (I5) If $t_1 + t_2 < N$, then $b_1 + b_2 = 0$; (I6) If $t_1 + t_2 + b_1 + b_2 \geq N$, then $t_1 + t_2 = N$; (I7) Once the cache is full, it remains full from that point onwards.

CAR maintains an adaptive parameter p , which it uses as a target for t_1 , the size of list T_1 . Consequently, $N - p$ is the target for t_2 . Using this guiding principle, it decides whether to evict an item from T_1 or T_2 in the event that a miss requires one of the pages to be replaced. The replacement policy can be summarized into two main points: (a) If the number of items in T_1 (barring the marked items at the head of the list) exceeds the target size, p , then remove a page from T_1 , else remove a page from T_2 or remove one of the marked pages at the head of the list; and (b) If $\ell_1 = t_1 + b_1 = N$, then remove a history page from B_1 , else remove a history page from B_2 . Since the details of the algorithm are complex, the actual pseudocode is provided (Figure 2 from [2]) in the Appendix.

4.1 Analyzing the Competitiveness of **CLOCK**

Before proving the competitiveness of **CAR**, since the algorithm relies on the simpler **CLOCK** algorithm, we first prove the competitiveness of **CLOCK**. We are not aware of the existence of a formal proof of competitiveness of **CLOCK**. Our result is formulated as the following theorem:

► **Theorem 7.** *Algorithm **CLOCK** is $2N$ -competitive.*

Proof. Let q be a page requested by the request sequence. We define $R[q]$ as follows. If q is in the cache maintained by **CLOCK**, but is not part of the cache maintained by **OPT**, then $R[q]$ is computed as follows. If q is unmarked, $R[q]$ equals the position of q counting counterclockwise from the tail. (If item q is at the tail, then $R[q] = 1$). If q is marked, then it equals N plus the position from the tail. Otherwise, $R[q] = 0$.

Following the proof for **ARC**, let C_{clock} and C_{opt} be the costs incurred by the algorithms **CLOCK** and **OPT**, and let $\sigma = \sigma_1, \sigma_2, \dots, \sigma_m$ be an arbitrary request sequence. As before, we partition σ into phases $P(0), P(1), \dots$, such that **CLOCK** has exactly N faults on $P(i)$, for every $i \geq 0$. Let the subsequence of the request sequence σ in phase $P(i)$ be $\sigma(i)$.

Let D be the set of pages that are in the cache maintained by **CLOCK**, but not in the cache maintained by **OPT**. We define the potential function as follows:

$$\Phi = \sum_{q \in D} R[q] \tag{4}$$

If both algorithms, **CLOCK** and **OPT**, start with the same cache contents, then phase $P(0)$ is easily handled since **OPT** will fault at least once and **CLOCK** faults exactly N times, making **CLOCK** N -competitive in this phase.

In order to analyze phase $P(i), i \geq 0$, we will prove that for every individual request, $\sigma_j \in \sigma(i)$:

$$C_{clock}(\sigma_j) + \Delta\Phi \leq 2N * C_{opt}(\sigma_j) \quad (5)$$

Summing up the above inequality for all requests, $\sigma_j \in \sigma(i)$ will prove inequality (5).

As before, we assume that request σ_j is processed in two distinct steps: first when **OPT** services the page request and, next when **CLOCK** services the request. We will show that inequality (5) is satisfied for each of the two steps.

When only **OPT** acts in this step, $C_{clock} = 0$, and t does not change. If **OPT** does not fault on this request, then $C_{OPT} = 0$. Since the contents of the cache maintained by **OPT** does not change, no pages get marked or unmarked, and the clock hand does not move. Thus, $\Delta\Phi = 0$.

If **OPT** faults on request σ_j , then $C_{OPT} = 1$. The contents of the cache maintained by **OPT** does change, which could affect the potential function. If a marked page in **OPT** gets evicted, since its R -value cannot exceed $2N$, the potential will change by at most $2N$. Thus, $C_{arc}(\sigma_j) + \Delta\Phi \leq 2N$, proving inequality (5).

Next we consider the steps when **CLOCK** services the request. If **CLOCK** has a hit, then there will be no change in potential since pages do not move around. If the requested page is marked by **CLOCK**, it does not change the potential either since the page is already in the cache maintained by **OPT** and is not considered for the potential function calculations. Finally, we consider the case when **CLOCK** has a miss. Since there are N pages in the cache maintained by **CLOCK**, at least one of those pages is guaranteed to be not part of the cache maintained by **OPT**. The newly requested page will be moved to the tail position. All other pages will move down by at least one position (in which case their R -value decreases by 1) or will lose their marked status (in which case their R -value decreases by a large amount, going down from $2N$ to 1). Either way, the decrease in potential will pay for **CLOCK**'s miss. This completes the proof of the theorem and the competitiveness analysis of the **CLOCK** algorithm. \blacktriangleleft

4.2 Analyzing the Competitiveness of CAR

Next, we analyze the competitiveness of **CAR**. The main result of this section is the following:

► **Theorem 8.** *Algorithm **CAR** is $21N$ -competitive.*

Proof. We introduce the following notation. Let q be any page in the request sequence. If q is in the cache maintained by **CAR**, then it is in one of the lists T_1, T_2, B_1, B_2 , and we define $P[q]$ to be the position of q in that list. For T_1 and T_2 , position is counted from the tail of the list to the head of the list in the counterclockwise direction. For B_1 and B_2 , position is counted from the MRU to the LRU. Recall that T_1^0 and T_2^0 are the unmarked pages in T_1 and T_2 , while T_1^1 and T_2^1 are the marked pages in T_1 and T_2 . Let D be the set of pages in $T_1 \cup T_2 \cup B_1 \cup B_2$, but not in the cache maintained by **OPT**. Let U be the set of pages contained in the intersection of $T_1 \cup T_2$ and the set of pages in the cache maintained by **OPT**.

As before, we associate each page with a value $R[q]$, which is defined as follows:

$$R[q] = \begin{cases} P[q], & \text{if } q \in B_1 \cup B_2 \\ 2 \cdot P[q] + b_1, & \text{if } q \in T_1^0 \\ 2 \cdot P[q] + b_2, & \text{if } q \in T_2^0 \\ 3N + 2 \cdot P[q] + b_1, & \text{if } q \in T_1^1 \\ 3N + 2 \cdot P[q] + b_2, & \text{if } q \in T_2^1 \\ 0, & \text{otherwise} \end{cases}$$

Finally, we define the potential function as follows:

$$\Phi = p + 2 \cdot (b_1 + t_1) - 3 \cdot |U| + 3 \cdot \sum_{q \in D} R[q] \quad (6)$$

We prove Theorem 8 by proving the following inequality for any phase $P(i), i \geq 0$.

$$C_{arc}(\sigma(i)) + \Delta\Phi \leq 21N \cdot C_{opt}(\sigma(i)) \quad (7)$$

Summing up the above inequality for all requests in all phases would prove the theorem. The arguments for phase $P(0)$ are identical to that for **ARC** and are not reproduced here. In order to analyze phase $P(i), i > 0$, we will prove that for every individual request, $\sigma_j \in \sigma(i)$:

$$C_{car}(\sigma_j) + \Delta\Phi \leq 21N \cdot C_{opt}(\sigma_j) \quad (8)$$

Summing up the above inequality for all requests, $\sigma_j \in \sigma(i)$ will prove inequality (7) for phase $P(i)$. As before, we assume that request σ_j is processed in two distinct steps: first when **OPT** services the page request and, next when **CAR** services the request. We will show that inequality (8) is satisfied for each of the two steps.

Step 1: OPT serves request σ_j .

Since only **OPT** acts in this step, $C_{CAR} = 0$, and t does not change. There are two possible cases: either **OPT** faults on σ_j or it does not.

If **OPT** does not fault on this request, then it is easy to see that $C_{OPT} = 0$ and $\Delta\Phi = 0$, thus satisfying inequality (8).

If **OPT** faults on request $\sigma(i)$, then $C_{OPT} = 1$. The contents of the cache maintained by **OPT** changes, with at most one page q whose $R[q]$ value changes. The maximal positive change will occur when the marked head page in T_2 is removed from **OPT**. Additionally, if $|T_2| + |B_2| = 2N$, then it could result in a potential change of $18N$. Since one element will be evicted from the cache maintained by **OPT**, $|U|$ may change by at most 1. Thus, the potential function Φ may change by at most $18N + 3 \leq 21N$. $C_{car}(\sigma_j) + \Delta\Phi \leq 21N$, proving inequality (8).

Step 2: CAR serves request σ_j .

As with the proof for **ARC**, there are four cases. While the structure is the same, the details are different. Case 1 deals with the case when **CAR** finds the page in its cache. The other three cases assume that **CAR** faults on this request because the item is not in $T_1 \cup T_2$. Cases 2 and 3 assume that the missing page is found recorded in the history in lists B_1 and B_2 , respectively. Case 4 assumes that the missing page is not recorded in history.

Case I: CAR has a page hit.

Clearly, the page was found in $T_1 \cup T_2$, and $C_{car} = 0$. Neither the cache nor the history lists maintained by **CAR** will change. Thus the contribution to $|U|$ does not change. Since

the page is not found in **CAR**'s history, $\Delta p = 0$. Since **OPT** has already served the page, the page is in **OPT**'s cache. Therefore, even if the page gets marked during this hit, its R -value does not contribute to the potential. Thus, $\Delta\Phi = 0$.

Next we will analyze the three cases when the requested page is not in **CAR**'s cache. Since $C_{car} = 1$, the change in potential must be at most -1 in order for inequality (8) to be satisfied. The following lemma is useful in understanding the potential change to the term $\sum_{q \in D} R[q]$.

► **Lemma 9.** *When **CAR** has a page miss, the change to the term $\sum_{q \in D} R[q]$ is at most 0.*

Proof. We examine the potential change based on the original location of the page(s) whose R -value changed.

Case $q \in B_1$ This implies that page q or some page in B_1 that was more recently used (higher R -value) was either discarded or requested (and consequently moved out). The contribution from the R -values of the pages in B_1 can therefore only go down or remain the same. If the page moved to T_2 , then it was because it was accessed, but then it is already in **OPT** and therefore does not contribute.

Case $q \in T_1^0$ A page in T_1^0 will decrease in R -value when b_1 decreases, or as a result of a call to REPLACE. This happens when some page is moved from T_1 to B_1 or from T_1 to T_2 (causing all pages to modify their position numbers and their R -values). However, these changes only decrease their R -values.

Case $q \in T_1^1$ The argument is exactly the same as for the previous case. In the event that the page itself is moved to T_2 it will lose its marked status which will cause a decrease in potential.

◀

The other three cases are omitted since they are similar to the above three. As with **ARC**'s analysis, we know that if **CAR** has a miss there is at least one page in either T_1 or T_2 that is not part of **OPT**. We proceed with analyzing the changes in the other terms in the potential function.

Case II: **CAR** has a miss and the missing page is in B_1

Note that in this case the value of p will change by at most +1, unless its value equals N , in which case it has no change. Thus $\Delta p \leq 1$. $\Delta R[q]$ will be a non-positive number. The reason why it is not guaranteed to be negative is because all the pages ordered above may be part of **OPT** and then REPLACE is called on T_1 . In this case $R[q]$ will not change. $\Delta|T_1| + |B_1|$ is guaranteed to decrease by at least one. **CAR** and **OPT** will still share the same pages.

$$\begin{aligned} \Delta\Phi &\leq 3 \cdot \Delta \sum_{q \in Car(history)/Opt} R[q] + 2 \cdot \Delta|B_1 + T_1| + \Delta p - 3 \cdot \Delta(Car \in Opt) \\ &\leq -1 \end{aligned}$$

Case III: **ARC** has a miss and the missing page is in B_2 .

Note that in this case the value of p will change by -1, if its value was positive, otherwise it has no change. Thus $\Delta p \leq 0$.

If we assume that REPLACE is called on T_2 and $R[q]$ doesn't then we could have

$$\begin{aligned} \Delta\Phi &\leq 3 \cdot \Delta \sum_{q \in D} R[q] + 2 \cdot \Delta(b_1 + t_1) + \Delta p - 3 \cdot \Delta|U| \\ &\leq 0 \end{aligned}$$

But this is not good enough since we need the potential change to be at most -1. When $\Delta p = -1$, then we get the required inequality $\Delta\Phi \leq -1$. Clearly, the difficulty is when $\Delta p = 0$.

Case III.1 If REPLACE is called on T_1 and a page in T_1 is not part of **OPT** then the contents of $R[q]$ will change by at least one. Thus $\Delta\Phi \leq -1$.

Case III.2 If REPLACE is called on T_1 and all pages in T_1 are part of **OPT** then the contents of $R[q]$ won't change from REPLACE. Since (lemma) $R[q]$ will change from B_2 moving to T_2 .

Case III.3 If REPLACE is called T_2 this means that $R[q]$ will definitely decrease...

Case IV: **CAR** has a miss and the missing page is not in $B_1 \cup B_2$

We consider three cases. First, when $\ell_1 = N$, **CAR** will evict the $LRU(B_1)$. Since there are N pages in L_1 at least one of the pages must not be part of **OPT**. By lemma() we know that the potential will change by a negative amount.

$$\begin{aligned} \Delta \Phi &\leq 3 \cdot \Delta \sum_{q \in D} R[q] + 2 \cdot \Delta(b_1 + t_1) - 3 \cdot \Delta|U| \\ &\leq -1 \end{aligned}$$

If $\ell_1 < N$ and $\ell = N$, then **CAR** will evict the $LRU(B_2)$. Since there are N pages in L_2 at least one of the pages must not be part of **OPT**. By lemma() we know that the potential will change by a negative amount.

$$\begin{aligned} \Delta \Phi &\leq 3 \cdot \Delta \sum_{q \in D} R[q] + 2 \cdot \Delta(b_1 + t_1) - 3 \cdot \Delta|U| \\ &\leq -1 \end{aligned}$$

The last case that will be examined will have no discards from history. Since no pages are discarded there is no guarantee that $R[q]$ will change. This also means that **CAR** and **OPT** will have one more page in common. Thus, the equation is

$$\begin{aligned} \Delta\Phi &\leq 3 \cdot \Delta \sum_{q \in D} R[q] + 2 \cdot \Delta(b_1 + t_1) - 3 \cdot \Delta|U| \\ &\leq -1 \end{aligned}$$

◀

5 Conclusions

Adaptive algorithms are tremendously important in situations where inputs are infinite on-line sequences and no single optimal algorithm exists for all inputs. Thus, different portions of the input sequence require different algorithms to provide optimal responses. Consequently, it is incumbent upon the algorithm to sense changes in the nature of the input sequence and adapt to these changes. Unfortunately, these algorithms are harder to analyze. We present an analysis of an important adaptive algorithm called **ARC** and show that it is competitive.

ACKNOWLEDGMENTS: This work was partly supported by NSF Grant (CNS-1018262) and the NSF Graduate Research Fellowship (DGE-1038321).

References

- 1 S. Albers. Online algorithms: a survey. *Mathematical Programming*, 97(1-2):3–26, 2003.
- 2 S. Bansal and D. S. Modha. CAR: CLOCK with adaptive replacement. In *Proceedings of the 3rd USENIX Conference on File and Storage Technologies*, FAST '04, pages 187–200, Berkeley, CA, USA, 2004. USENIX Association.
- 3 F. J. Corbato. A paging experiment with the MULTICS system. Technical report, DTIC Document, 1968.
- 4 M. B. Friedman. Windows NT page replacement policies. In *Proceedings of the Intl. CMG Conference*, pages 234–244, 1999.
- 5 D. S. Hochbaum, editor. *Approximation algorithms for NP-hard problems*. PWS Publishing Co., Boston, MA, USA, 1997.
- 6 A. Janapsatya, A. Ignjatovic, J. Peddersen, and S. Parameswaran. Dueling CLOCK: adaptive cache replacement policy based on the CLOCK algorithm. In *Design, Automation & Test in Europe Conference & Exhibition (DATE), 2010*, pages 920–925. IEEE, 2010.
- 7 S. Jiang, F. Chen, and X. Zhang. CLOCK-Pro: An effective improvement of the CLOCK replacement. In *USENIX Annual Technical Conference, General Track*, pages 323–336, 2005.
- 8 S. Jiang and X. Zhang. LIRS: An efficient low inter-reference recency set replacement policy to improve buffer cache performance. In *Proc. ACM Sigmetrics Conf.*, pages 297–306. ACM Press, 2002.
- 9 T. Johnson and D. Shasha. 2Q: A low overhead high performance buffer management replacement algorithm. In *Proc. of VLDB*, pages 297–306, 1994.
- 10 D. Lee, J. Choi, J. H. Kim, S. H. Noh, S. L. Min, Y. Cho, and C. S. Kim. LRFU: A spectrum of policies that subsumes the least recently used and least frequently used policies. *IEEE Trans. Comput.*, 50(12):1352–1361, December 2001.
- 11 N. Megiddo and D. S. Modha. ARC: A self-tuning, low overhead replacement cache. In *Proceedings of the 2nd USENIX Conference on File and Storage Technologies*, FAST '03, pages 115–130, Berkeley, CA, USA, 2003. USENIX Association.
- 12 N. Megiddo and D. S. Modha. Outperforming LRU with an adaptive replacement cache algorithm. *IEEE Computer*, 37(4):58–65, 2004.
- 13 E. J. O’Neil, P. E. O’Neil, and G. Weikum. The LRU-K page replacement algorithm for database disk buffering. *SIGMOD Rec.*, 22(2):297–306, June 1993.
- 14 D. D. Sleator and R. E. Tarjan. Amortized efficiency of list update and paging rules. *Commun. ACM*, 28(2):202–208, February 1985.

6 Appendix

Pseudocode: ARC(x)INPUT: The requested page x INITIALIZATION: Set $p = 0$ and set lists T_1 , B_1 , T_2 , and B_2 to empty

```

1: if ( $x$  is in  $T_1 \cup T_2$ ) then                                     ▷ cache hit
2:   Move  $x$  to the top of  $T_2$ .
3: else if ( $x$  is in  $B_1$ ) then
4:   ADAPTATION: Update  $p = \min\{p + \max\{1, |B_2|/|B_1|\}, N\}$ 
5:   REPLACE()
6:   Move  $x$  to the top of  $T_2$  and place it in cache.
7: else if ( $x$  is in  $B_2$ ) then
8:   ADAPTATION: Update:  $p = \max\{p - \max\{1, |B_1|/|B_2|\}, 0\}$ 
9:   REPLACE()
10:  Move  $x$  to the top of  $T_2$  and place it in cache.
11: else                                                           ▷ cache and directory miss
12:  if ( $|T_1| + |B_1| = N$ ) then
13:    if ( $|T_1| < N$ ) then
14:      Discard LRU page in  $B_1$ .
15:      REPLACE()
16:    else
17:      Discard LRU page in  $T_1$  and remove it from cache.
18:    end if
19:  else if ( $(|T_1| + |B_1| < N)$  and  $(|T_1| + |T_2| + |B_1| + |B_2| \geq N)$ ) then
20:    if ( $|T_1| + |T_2| + |B_1| + |B_2| = 2N$ ) then
21:      Discard LRU page in  $B_2$ .
22:    end if
23:    REPLACE()
24:  end if
25: end if

```

REPLACE()

```

26: if ( $(|T_1| \geq 1)$  and  $((x$  is in  $B_2$  and  $|T_1| = p)$  or  $(|T_1| > p))$ ) then
27:   Delete LRU page in  $T_1$  (also remove it from cache); move it to MRU position in  $B_1$ .
28: else
29:   Delete LRU page in  $T_2$  (also remove it from cache); move it to MRU position in  $B_2$ .
30: end if

```

Pseudocode: CAR(x)

 INPUT: The requested page x

 INITIALIZATION: Set $p = 0$ and set lists T_1 , B_1 , T_2 , and B_2 to empty

```

1: if ( $x$  is in  $T_1 \cup T_2$ ) then                                     ▷ cache hit
2:   Set the page reference bit for  $x$  to one.
3: else                                                           ▷ cache miss
4:   if ( $|T_1| + |T_2| = N$ ) then                                     ▷ cache full, replace a page from cache
5:     REPLACE()                                                  ▷ cache directory replacement
6:     if ( $(x$  is not in  $B_1 \cup B_2$ ) and  $(|T_1| + |B_1| = N)$ ) then
7:       Discard LRU page in  $B_1$ .
8:     else if ( $(|T_1| + |T_2| + |B_1| + |B_2| = 2N)$  and  $(x$  is not in  $B_1 \cup B_2)$ ) then
9:       Discard LRU page in  $B_2$ .
10:    end if
11:  end if
12:  if ( $x$  is not in  $B_1 \cup B_2$ ) then                             ▷ cache directory miss
13:    Insert  $x$  at the tail of  $T_1$ . Set the page reference bit of  $x$  to 0.
14:  else if ( $x$  is in  $B_1$ ) then                                   ▷ cache directory hit
15:    ADAPTATION: Update  $p = \min\{p + \max\{1, |B_2|/|B_1|\}, N\}$ 
16:    Move  $x$  at the tail of  $T_2$ . Set the page reference bit of  $x$  to 0.
17:  else                                                         ▷ cache directory hit
18:    ADAPTATION: Update:  $p = \max\{p - \max\{1, |B_1|/|B_2|\}, 0\}$ 
19:    Move  $x$  at the tail of  $T_2$ . Set the page reference bit of  $x$  to 0.
20:  end if
21: end if

```

 REPLACE()

```

22: found = 0
23: repeat
24:   if ( $|T_1| \geq \max(1, p)$ ) then
25:     if (page reference bit of head page in  $T_1$  is 0) then
26:       found = 1
27:       Demote head page in  $T_1$  and make it MRU page in  $B_1$ 
28:     else
29:       Set the page reference bit of head page in  $T_1$  to 0, and make it tail page in  $T_2$ .
30:     end if
31:   else
32:     if (page reference bit of head page in  $T_2$  is 0) then
33:       found = 1
34:       Demote head page in  $T_2$  and make it MRU page in  $B_2$ .
35:     else
36:       Set page reference bit of head page in  $T_2$  to 0, and make it tail page in  $T_2$ .
37:     end if
38:   end if
39: until (found)

```
